

**APPLICATION
FOR
UNITED STATES
LETTERS
PATENT**

**METHOD AND APPARATUS FOR
RECEIVING DATA**

INVENTORS:

DUNNING, DAVID S.

ABHAYAGUNAWARDHANA, CHAMATH

DROTTAR, KENNETH

JENSEN, RICHARD

ASSIGNEE: INTEL CORPORATION

Intel Reference Number P9455

FILED: December 27, 2000

METHOD AND APPARATUS FOR RECEIVING DATA

FIELD OF THE INVENTION

The invention relates generally to methods and apparatuses for receiving
5 data, and more particularly to a method and apparatus for receiving data over a
noisy channel, such as that over which servers communicate.

BACKGROUND OF THE INVENTION

In serial data communication, the channel through which information is
10 transmitted imposes a limit on the bandwidth capacity, or rate at which information
may pass through the channel. One significant limitation on achievable bandwidth
is known as inter-symbol interference, which occurs when a portion of a signal
representative of one bit of information interferes with a different portion of the
signal representative of a different bit of information. While the portions of the
15 signal corresponding to respective bits are transmitted such that they occur at
different times (i.e., within respective bit periods), efforts to exploit more and more
of a media's theoretical available bandwidth results in a concomitant increase in the
likelihood that a transmitted pulse will be "chirped" or spread, at lower power, over
consecutive bit periods before reaching the receiver.

20 Besides attenuating the maximum energy contained in any one-bit period,
the affected pulses interfere with adjacent bits. Since adjacent bits may not be
correlated consistently, the interfering bits or pulses have a direct and substantial
impact on the effective noise power and thus reduces the signal to noise ratio.

While it is known to employ equalization in the receiver using analog
25 circuitry, such techniques are generally considered to be too complex and or
expensive for applications involving transmission at higher bit rates.

The invention is therefore directed to the problem of developing a system
and method capable of extending the benefits of equalization to higher bit rate
transmission.

30

BRIEF DESCRIPTION OF THE DRAWINGS

FIG 1 illustrates a block diagram of an illustrative high-speed data transmission system implementing digital equalization according to the invention.

FIG 2 depicts an exemplary embodiment of an edge-based receiver
5 implementing digital equalization according to the invention.

FIG 3 depicts an exemplary embodiment of an edge processor, which forms part of the exemplary edge-based receiver of FIG 2 and performs digital equalization to compensate for deterministic jitter of data in a received data signal.

FIG 4 depicts the bias introduced in the received signal by deterministic
10 jitter.

FIG 5 depicts the possible regions where a transition can be registered and the associated decision, for an illustrative 2.5 Gbps receiver implementation employing four clock phases per bit time.

FIG 6 depicts an exemplary embodiment of the using step in the
15 embodiment of FIG 7 according to another aspect of the invention.

FIG 7 depicts an exemplary embodiment of the final two steps of the embodiment of FIG 6 according to another aspect of the invention.

DETAILED DESCRIPTION

20 The invention solves these and other problems by a method and apparatus for providing equalization for a communication channel which avoids the complexity and additional components required by prior art equalization techniques. By avoiding such complexity, the invention supports high data rates.

It is worthy to note that any reference herein to “one embodiment” or “an
25 embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment.

In accordance with an illustrative embodiment of the invention, edge
30 transition samples, along with a history of data transmitted over the communication channel, are monitored and given weight in a determination on the amount of equalization to be applied.

A method of recovering data from a modulated data signal in accordance with an illustrative embodiment of the invention includes tracking a transmitted clock with a plurality of locally-generated clock phases, estimating an average phase of previously detected edges, registering a pulse edge in the received stream of data at a transition phase corresponding to one of the plurality of locally-generated clock phases, determining whether a first symbol was received multiple times consecutively prior to the registered pulse edge, and using the determination of whether the first symbol was received multiple times consecutively to determine whether to assign a registered pulse edge to a current or subsequent bit period.

System Overview

FIG 1 is a block diagram of an exemplary embodiment of a high-speed communications system 10 according to one aspect of the invention. The high-speed communication system 10 of the illustrative embodiment includes a transmitter 12 for producing a differentially encoded signal. To this end, the transmitter contains a multiple phase reference clock 14 operatively associated with a phase locked loop (PLL) 16 for generating a plurality of clock phases synchronized with reference clock 14. Illustrative system 10 further includes a receiver 20, which is coupled to the transmitter via a suitable media 21, which defines a data channel.

One possible use for the exemplary high-speed communication system 10 includes communications between servers, which are interconnected via a bus or network. These servers need to transmit data at extremely high speeds and with few data errors. Consequently, the edge-based receiver of the invention can be used to further improve the data rate of such inter-server communications.

An exemplary embodiment of a computer system according to one aspect of the invention includes several modular components communicating with each other. Each of the modular components includes a transmitter and receiver, or alternatively, a transceiver. The transmitter transmits a modulated data signal that is subsequently received by the receiver.

Turning to FIG 6, and as will be described in more detail, the receiver tracks a transmitted clock using multiple locally-generated clock phases (step 71). The

receiver then estimates an average phase of the previously detected edges (step 72) and registers a pulse edge in the received stream of data at a transition phase corresponding to one of the locally-generated clock phases (step 73). Next, the receiver determines whether a first symbol was received multiple times
5 consecutively prior to the detected pulse edge (step 74) and uses this determination of whether the first symbol was received multiple times consecutively in a receiver decision process (step 75).

The transmitter, receiver and/or transceiver functions above may be performed using software. As such, these instructions may be encoded on
10 computer-readable media and subsequently executed by a processor. Any computer readable media may suffice, such as magnetic disks, optical disks, CD-ROMs, flash memory, EEPROMs, EPROMs, bubble memory, etc., to name only a few examples.

15 **Receiver Description**

The receiver employs a multi-phase clock generator that runs plesiochronously with respect to the transmit clock. A frequency difference between the two clock bases is specified to be less than a few hundreds parts per millions. The receiver functionality is divided into three main sections, i.e., an edge
20 detector section, an edge processing section and an elastic buffer section. A high-level block diagram of this illustrative structure is shown in FIG 2, which depicts edge buffer 22, edge processing section 24, and multi-phase clock generator 26.

Essentially, edge buffer 22 receives the coded transmission signal, which, in the illustrative communication system depicted in FIG 1, is a differentially encoded
25 signal. The edge signals produced by edge buffer 22 are forwarded to edge processor 24 upon receipt of an enable signal from edge processor 24. The edge processor 24 sends the data and a recovered clock phase ϕ_z to elastic buffer 26, out of which the data is clocked. Multi-phase clock 28 provides references to edge processor 24 for use in determining the recovered clock phase ϕ_z . In accordance
30 with an illustrative embodiment of the invention, and in a manner to be described in greater detail later, data history accumulated over a number of preceding, consecutive bit periods by the elastic buffer 26 is used to assign each pulse edge or

transition, as it is detected, to either a current or a subsequent bit period.

A receiver constructed in accordance with an illustrative embodiment of the invention operates by detecting pulse edges or transitions, also referred to as “zero crossings”, that occur in the input transmitted waveform or signal. As used herein, the term “zero crossings” or edges is used to refer to the start of a new symbol or bit, such as the time instances when two differential input signals cross each other. For a bit encoded using only one signal, the edges or zero crossings refer to the start of the symbol, such as the rising or falling edge of a pulse or turning on and off of a light source. Therefore, an edge or zero crossing signals a change in the transmitted bit as well as the absence of any data changes since the last edge.

The invention, through application of observed phase noise behavior across transmission lines, avoids the penalty of quantization jitter in systems, such as the illustrative differential encoding/decoding system described above, using a digital equalization technique. This digital equalization is performed, in accordance with the illustrative embodiment, by the edge processor 24, in a manner that will now be described.

Digital Equalization

As exemplified by an illustrative embodiment of the invention depicted in FIG 3, digital equalization is facilitated by a data recovery unit 30 which forms a part of the edge processing unit 24. The purpose of data recovery unit 30 is to establish the final data output, by using sampled information received from a synchronizer and alignment unit 32, and an estimated mean zero crossing phase provided by tracking clock unit 34. In this process, the data recovery unit will remove the deterministic jitter associated with the edges, thereby relaxing the timing margins at the receiver. The samples received indicate the received edges synchronized to a multi-phase local clock. In effect, the multiple phases are used to bin the edge instances (time of occurrence) into several time intervals.

The data recovery unit uses the recovered data history to subtract the deterministic jitter associated with the edges. In order to subtract out the deterministic jitter (DJ) component, it is first necessary to estimate the composition of random jitter (RJ) and deterministic jitter seen at the input to the synchronizers.

For purposes of illustration, it is sufficient to estimate the worst-case jitter seen at the input to the synchronizer to account for the jitter added by the front-end of the receiver. In the discussion that follows, exemplary jitter estimates and associated assumptions are discussed, followed by a description of an illustrative decision mechanism which may be employed to carry out the objectives of the invention.

Table 1 shows the jitter budget used in an exemplary embodiment of a communication system at an illustrative data rate of 2.5 Gb/s, and indicates a 0.35 unit interval (UI) eye opening at the receiver circuit input.

Table 1: Interconnect Budget for 2.5 Gbps Operation.

	DJ (UI) pk-pk	RJ (UI) pk-pk	Total (UI) pk-pk
Transmitter	0.17	0.18	0.35
Channel	0.24	0.06	0.30
Total	0.41	0.24	0.65

The exemplary receiver is configured to work with an eye opening of 0.15 UI at the input to the synchronizers, at 2.5 Gbps. For purposes of the illustrative communication system depicted in FIG 1, it is appropriate to consider that the receiver produces the same random jitter as the transmitter, i.e., 0.18 UI since the dominant random jitter component is presumed to be the phase locked loop (PLL), an element common to both transmitter and receiver. Deterministic jitter (DJ) of 0.02 UI is allocated to the receiver, leaving an eye opening of 0.15 UI. For purposes of analyzing the timing requirements, it is convenient to allocate all the received jitter components to the transmitter side, and assume zero jitter in the receiver. This approach will yield the following jitter numbers at the input to the synchronizer.

Total RJ seen at the synchronizer inputs = 0.18 UI + 0.18 UI + 0.06 UI = 0.42 UI

Total DJ seen at the synchronizer inputs = 1.00 UI - 0.15 UI - 0.42 UI = 0.43 UI

Table 2: Jitter Budget at Input to Synchronizers for 2.5 Gbps Operation.

	DJ (UI) pk-pk	RJ (UI) pk-pk	Total (UI) pk-pk
Input to the synchronizers	0.43	0.42	0.85
Receiver margin			0.15

The total DJ at a given bit time is related to the history of the received data. This leads to the assumption that the worst case DJ of 0.43 UI seen at the synchronizers corresponds to a lone transition ('000001' or '111110') or a trailing edge of a lone pulse ('0000010' or '1111101').

In order to remove the deterministic jitter, it is proposed to quantify the deterministic jitter by looking at the data history. Here again, for ease of explanation, analysis continues on the basis of one-sided jitter values, from peak to peak quantities, in order to quantify the data dependent effects. For purposes of illustrative application of the underlying inventive principles to an exemplary receiver, it is assumed that deterministic jitter increases linearly with the number of consecutive symbols. While a more accurate model may be obtained by using specific characteristics of the actual channel, the foregoing generalization is sufficient for illustrative purposes. While it might further be noted that deterministic jitter also depends upon the cable length, such effects may be ignored since the analysis here proceeds on a worst-case basis. It suffices to say that shorter cables will produce larger eye openings, and hence a higher signal to noise ratio.

Table 3: Magnitude of the DJ with Number of Consecutive Symbols

# consecutive symbols	DJ (UI) one-sided
5	0.22
4	0.16
3	0.11
2	0.06
0	0

The amount of deterministic jitter present in the edges tends to move them in opposite directions depending on the history of the data. If the current edge

corresponds to a lone transition, the deterministic jitter always tends to delay it in time, i.e., *late edge*. On the other hand, if the current edge corresponds to the trailing edge of a lone pulse, then deterministic jitter has the opposite effect, advancing the edge in time, i.e., *early edge*. The bias introduced by the

5 deterministic jitter is shown in FIG 4. By quantifying the deterministic effects of the lone transition, trailing edge of the lone pulse and their *friends* (the patterns ‘00001’, ‘0001’, ‘001’ and their inversions are considered to be friends of the lone transition, and the patterns ‘000010’, ‘00010’, ‘0010’ and their inversions are considered to be “friends” of the lone pulse), it is possible to effectively remove

10 deterministic jitter from the samples.

The above reasoning can be further extended by considering the friends of the lone transitions and the trailing edge of the lone pulse. The edges, which are similar to the lone transition and the trailing edge of the lone pulse, i.e., their friends, have similar effects on the transition, at a lesser magnitude. For example,

15 the pattern ‘00001’ will have considerable deterministic jitter to delay the transition. This is also shown in FIG 4.

Decision Mechanism

In accordance with an illustrative embodiment of a receiver, the decision

20 mechanism employs three pieces of information in order to recover the data.

1. Currently selected mean zero-crossing phase ϕ_z .
 2. The clock phase ϕ_t (transition phase) which registered the transition, if any.
 3. The history of the recovered data.
- 25 Data recovery unit (DRU) 30 assumes the received bit periods to start at ϕ_z . If a transition is registered with ϕ_t , then data recovery unit 30 must toggle the data output, in order to reflect the transition. However, it is first necessary to determine the bit period to which the transition belongs. Depending on t , there are two possibilities. It may belong to the current bit period or the next bit period. FIG 5
- 30 shows the possible regions where a transition can be registered and the associated decision, for an illustrative 2.5 Gbps implementation with an exemplary four clock phases per bit time.

According to FIG 5, the current mean zero crossing phase $\varphi_Z = \varphi_3$. There are three regions (or three possible values for the transition phase φ_t) where the decision is fairly straightforward. Since there is a guaranteed eye opening at the center, it can be concluded that any transition registered to the left of the eye opening correspond to the current bit period and all transitions to the right of the eye correspond to the next bit period.

In the example, a transition registered with φ_3 or φ_0 indicates that it belongs to the current bit period, whereas a transition registered with φ_Z indicates that it belongs to the next bit period. However, if a transition is registered with φ_1 , the decision is ambiguous. That is, it might either correspond to a late transition from the current bit time or to an early transition from the next bit time. So long as the eye opening is smaller than the spacing between two adjacent clock phases, this ambiguity will exist. The region where the decision is ambiguous shall be referred to herein as the "dead zone."

It is possible to adopt several methods to avoid the decision ambiguity when a transition happens in the dead zone. One possible method is to increase the number of clock phases per bit time. However, it is difficult to generate and route high-speed, multi-phase clocks with low jitter. Therefore, this option may not be suitable beyond a certain number of clock phases, and also does not scale with increasing jitter.

Alternatively, the history of the data may be used to resolve the ambiguity, effectively removing the deterministic jitter associated with the transition.

A third method is to separate out the transitions into two categories, i.e., positive and negative zero crossings, and assign a polarity to the transition. By distinguishing between the polarities, an extra rule can be added to the decision process, i.e., consecutive transitions must always have different polarities. In accordance with an illustrative embodiment of receiver, both the second and third approaches are employed.

First, the deterministic jitter removal is achieved using data history. With continued reference to FIG. 5, the transition regions shown therein depict a gap between a late edge of the current bit period and an early edge belonging to the next bit period. It should be noted that the data is based on the jitter specifications and

the assumption that deterministic jitter scales linearly. The illustrative data implies that only two extreme transition cases may fall in the dead zone, i.e., lone transition with five consecutive symbols and the trailing edge of the lone pulse with five consecutive symbols. These two cases can be easily distinguished due to the fact
5 that deterministic jitter gets added in opposite directions.

Also notice that the detection of the trailing edge of the lone pulse is easier than the detection of the lone transition, in the sense that we only need to detect a change in the last two bits. So the information obtained from deterministic jitter removal method may be simplified as follows: (1) if the last two bits are different
10 AND there is a dead zone transition, the registered transition or edge is assigned to the next bit period and (2) if the last two bits are equal AND there is a dead zone transition, the registered transition or edge is assigned to the current bit period.

The decision can be further enhanced by including the mean zero crossing location information provided by tracking clock unit 34 (FIG 3). Such information
15 provides greater accuracy in assigning the transition ranges established in the diagram, by reducing the error between the true mean zero crossing and the estimated mean zero crossing. The worst-case error in the phase estimation assumed in connection with the illustrative representation of FIG. 5 is ± 0.0625 UI.

Previous digital implementations of equalization require taking samples at a
20 rate considerably faster than the actual signal. If the signal is produced anywhere near the physical signaling limits of the silicon, analog circuitry must be used to recover the data. Analog circuitry is large, power consuming and does not replicate well. Receivers embodying the invention can recover data at the high rate of speed normally associated with analog circuitry while taking advantage of the power
25 savings, silicon substrate area savings and design/algorithmic repeatability.

The invention can be illustrated through a truth table and a circuit implementation. In the truth table, the data is extracted from four samples during any given bit period (UI). Digital averaging might further improve the results obtained by centering all marginal (or dead zone) transitions to the center of each
30 sample window. Again, these samples are based on several observations about the nature of phase distortions (or noise) across a transmission line, that is,

deterministic jitter distributes in a linear fashion per bit order and bit patterns bias average zero crossings late (dead zone transitions early).

FIG 6 depicts an exemplary embodiment of a method 60 according to one aspect of the invention. According to this embodiment, the transmitted clock is tracked with multiple locally-generated clock phases (step 71). The average phase of the detected edges is estimated (step 72). A pulse edge in the data stream is registered at the transition phase corresponding to one of the multiple locally-generated clock phases (step 73). Next, it is determined whether the first symbol was received multiple times consecutively (step 74). This determination is then used in the receiver decision process (step 75).

Turning to FIG 7, an exemplary embodiment of alternative embodiments of the interaction of the determining step 74 and the using step 75. In one embodiment, the process determines whether the registered pulse from step 73 is a lone transition (step 81). If not, then the process ends without a decision. If the registered pulse is a lone transition, then the process determines whether the lone transition has been preceded by a number of symbols above a predetermined threshold (step 82). If the threshold has not been exceeded, then the process determines whether the transition occurs before the eye opening in the received data signal (step 83). If so, the process determines that the registered pulse is associated with the current bit period (step 84). If not, the process determines that the registered pulse is associated with the next bit period (step 85). Returning back to step 82, if the threshold of consecutive symbols has been exceeded, the process determines whether the eye opening lies between the transition phase and the next phase (step 86). If so, the registered pulse edge is associated with the current bit period (step 84). If the eye opening does not lie between the transition phase and the next phase, then the registered pulse may belong to the next bit period or not.

The above process can be repeated based on determining whether the registered pulse is a trailing edge or not (step 81). The remaining steps (82-86) repeat as shown.

Although various embodiments are specifically illustrated and described herein, it will be appreciated that modifications and variations of the invention are covered by the above teachings and within the purview of the appended claims

without departing from the spirit and intended scope of the invention. For example, while several of the embodiments depict the use of four clock phases, other numbers (n) of clock phases will suffice, such as only two or more than four. In addition, while some of the above embodiments use a counter and a maximizer to
5 determine a moving average of the phase of the edges, any technique for calculating the moving average will suffice. Furthermore, while some of the above
embodiments use differential coding, any coding scheme, or no coding, will suffice. These examples should not be interpreted to limit the modifications and variations
of the invention covered by the claims but are merely illustrative of possible
10 variations.